

A SUPERWISE® LEARN KIT · OCTOBER 2026

AI terms that matter: the words in your contracts and policies

Twenty-seven terms a small or mid-size firm meets in AI vendor contracts, privacy terms, and its own policies. Each one in plain English, with the question to ask and where it shows up in the paperwork.

How to use this glossary

Vendor terms use ordinary words with narrow meanings. “Zero data retention” names a product setting with a list of exceptions, and “data residency” can mean storage in one country and processing in another. Each entry here gives the meaning a non-technical reader can repeat to a colleague, and the one question that tests it.

- 01 Read the term where you found it.** The same word can mean different things in an order form, a data processing agreement, and a help-center page. Part 5 shows which documents usually carry which terms.
- 02 Ask the question under the entry.** Put it to the vendor in writing, and keep the answer with the contract.
- 03 Check the date.** Vendor examples carry the date they were read. Help-center pages change without notice, so the signed agreement and its effective date decide.

PART	TERMS
2 • Your data	Zero data retention, retention period, training on your data, data residency, encryption, de-identification
3 • Contracts and parties	Business associate agreement, data processing agreement, controller, processor, subprocessor, incident and notice clock, SOC 2 report
4 • AI risks and controls	Prompt injection, hallucination, retrieval-augmented generation, guardrail, AI gateway, data loss prevention, shadow AI, human oversight
5 • Your own paperwork	Acceptable use policy, AI inventory, impact assessment, audit log, model card, system card, and a map of where each term lives

PLAIN ENGLISH • TERM OF ART

An everyday word that a contract, a law, or a vendor gives a specific meaning. The specific meaning is the one that binds, and it is usually defined near the front of the document or in a linked page.

Concepts live on the website. For how AI works (model weights, tokens, training, inference), see the free glossary at superwise.ai/learn/glossary/. This kit covers the words that decide what a vendor has promised you.

Not legal advice. Several of these terms have no single binding meaning across vendors and jurisdictions. The contract, the applicable law, and your counsel decide what one means for your firm.

Where it goes, and how long it stays

ZERO DATA RETENTION (ZDR)

An arrangement under which the vendor keeps no copy of your prompts and outputs once a request is answered, for the products it names. It usually needs the vendor's approval. OpenAI grants it on approval and lists API endpoints that can still store data with it on [1]. Anthropic applies it to eligible API products and Claude Code for Enterprise, and still keeps safety classifier results [2].

Ask: Which products and endpoints qualify, what is still kept, and where is it in our contract?

RETENTION PERIOD

How long a vendor keeps each kind of data, such as chats, files, logs, and safety records, before deleting it. Each kind can have its own clock. OpenAI keeps API abuse-monitoring logs for up to 30 days by default [1]. Anthropic keeps prompts and outputs for its designated covered models for 30 days [3].

Ask: How long is each kind of data kept, and can an admin shorten it?

TRAINING ON YOUR DATA

Contract wording on whether your prompts, outputs, files, or feedback may be used to train or improve the vendor's models. OpenAI keeps business data out of training by default unless you opt in [4]. Anthropic's commercial terms make exceptions for feedback you report and programs you choose to join [5].

Ask: Is it off by default, and can an admin stop users from opting in?

DATA RESIDENCY

The region where your data is stored, and sometimes where it is processed. Storage, processing, and support access can each sit in a different place. OpenAI stores ChatGPT Enterprise content at rest in 10 regions and offers in-region GPU inference in three; logins, billing, and workspace metadata can be held elsewhere [6].

Ask: Which data stays in the region: storage, processing, logs, backups, and support access?

ENCRYPTION AT REST AND IN TRANSIT

Scrambling data while it is stored and while it moves, so only someone holding the key can read it. An AI tool has to read content to work on it. Google states that Gemini in Workspace cannot act on email and files encrypted with customer-controlled keys [7].

Ask: Where is the data decrypted, and who can use the key?

DE-IDENTIFICATION

Removing the link between information and the person it describes. For health information, HIPAA recognizes two methods: removing 18 listed identifiers (Safe Harbor), or a qualified expert's documented determination [8]. Removing a name alone meets neither.

Ask: Which method was used, and who checked the result?

Who promises what, to whom

BUSINESS ASSOCIATE AGREEMENT (BAA)

The contract HIPAA requires before a vendor creates, receives, keeps, or sends protected health information for a healthcare practice or health plan. It sets what the vendor may do with the information and how it must protect it [9]. AI vendors offer one only for named plans and features [10, 11].

Ask: Will you sign a BAA for this exact plan, and which features does it exclude?

DATA PROCESSING AGREEMENT (DPA)

The contract that governs how a vendor handles personal data for you: what it processes, for how long, on whose instructions, with what security, and what happens to the data at the end. Article 28 of the EU's General Data Protection Regulation lists what it must cover [12].

Ask: Is the DPA part of our order, and does it cover every product and region we use?

CONTROLLER

The organization that decides why and how personal data is used [12]. For your clients' data, that is usually your firm.

Ask: Does the vendor ever use our data for its own purposes, which would make it a controller too?

PROCESSOR

An organization that handles personal data on a controller's behalf and on its instructions [12]. An AI vendor serving your firm is usually a processor.

Ask: Does the vendor act only on our documented instructions?

SUBPROCESSOR

A company the vendor engages to process your data, such as a cloud host or a model provider. Under GDPR the vendor needs the controller's authorization to add one, and the same data protection duties pass down to it [12].

Ask: Where is the current list, how are changes announced, and can we object?

INCIDENT AND NOTICE CLOCK

An incident is an event that harms, or could harm, the confidentiality, integrity, or availability of data or a system. Each contract and law sets its own definition and deadline. HIPAA gives a practice no more than 60 days from discovering a breach to notify the people affected [13].

Ask: How fast will the vendor tell us, and is that fast enough for our own deadlines?

SOC 2 REPORT

An independent CPA's examination of a vendor's controls relevant to security, availability, processing integrity, confidentiality, or privacy [14]. It covers the system and the period it describes, and vendors usually share it under a nondisclosure agreement.

Ask: May we see the report, and does it cover the product we are buying?

How AI tools go wrong, and what sits around them

PROMPT INJECTION

An attack that hides instructions in content the AI reads, so it follows the attacker. NIST defines it as exploiting the joining of untrusted input with a prompt built by a higher-trust party [15]. It can arrive directly from a user, or indirectly through a document, web page, or email the tool reads [16].

Ask: Which actions need a person's approval, and how are hidden instructions in documents tested?

HALLUCINATION (CONFABULATION)

Content an AI states confidently that is false. NIST calls it confabulation and lists it among the risks of generative AI to measure and manage [17].

Ask: What error rate was measured, on what test, and who checks high-stakes answers?

RETRIEVAL-AUGMENTED GENERATION (RAG)

A design in which the tool looks up documents and hands them to the model as context for its answer [17]. It ties answers to your sources. Whether the answer is right, and whether the user may see each document, are separate questions.

Ask: Which sources can it search, and does it respect each user's access?

GUARDRAIL

A technical or procedural constraint that detects, flags, filters, blocks, or routes for review an unwanted AI input, output, or action [17].

Ask: What does it check, and what happens when it triggers?

AI GATEWAY

A control point between people or applications and the AI models they call, where rules apply to each request that passes through it. An LLM gateway is the same idea for language-model traffic [18].

Ask: Which tools and which traffic pass through it?

DATA LOSS PREVENTION (DLP)

Policies and tools that find sensitive data and warn about, block, or record its movement out of the organization [19]. AI DLP applies the same idea to prompts and uploads.

Ask: Which channels and data types are covered, and how are false alarms handled?

SHADOW AI

AI tools, accounts, or agents used for work without the organization's knowledge or approval. Microsoft describes it as consumer AI apps and agents running without IT visibility or approval [20].

Ask: How would we know which AI tools staff use, and on which devices?

HUMAN OVERSIGHT

Named people with the authority and the information to monitor an AI system, step in, override it, or stop it, as the risk requires. The EU AI Act sets this out for high-risk systems in Article 14 [21].

Ask: Who reviews which decisions, and can they stop the action?

The documents your firm keeps

ACCEPTABLE USE POLICY

The firm's written rules for AI at work: which tools staff may use, which information stays out, who checks the output, and how to report a problem. A template is in the *AI acceptable use policy* kit [22].

Ask: Who owns it, and when was it last reviewed?

AI INVENTORY

A maintained list of the AI tools and uses in the firm, each with an owner, a purpose, the data involved, the vendor, and its status. The NIST AI Risk Management Framework places it among its governance practices [23].

Ask: Does it include AI features inside the software we already use?

IMPACT ASSESSMENT

A written review of what a system is for, who it affects, its risks and controls, and the risk that remains, done before use and repeated when something changes [23].

Ask: What triggers a new review, and who signs it off?

AUDIT LOG

A time-ordered record of who did what, when, and in which system, kept for accountability and investigation [19]. Its value depends on what it covers and how long it lasts. OpenAI's Compliance Logs Platform keeps logs for 30 days, so a firm that wants longer has to download them [24].

Ask: Can we export it, and how long is it kept?

MODEL CARD

A short document describing a model's intended uses, how it was evaluated, and its known limits [25].

Ask: Is it current for the exact model version we use?

SYSTEM CARD

A vendor document for a deployed AI product, covering its evaluations, safeguards, and limits, which can reach beyond the model underneath. No standard fixes its contents, so two vendors' system cards can cover different ground.

Ask: Does it describe the product as we will configure it?

Where each term lives

DOCUMENT	TERMS YOU WILL FIND THERE	WHAT TO CHECK
Order form or master agreement	Retention period, training on your data, incident, subprocessor	It names your exact plan and links the documents below.
Data processing agreement	Controller, processor, subprocessor, data residency, deletion	It covers every product and region you use.
Business associate agreement	Protected health information, permitted uses, breach notice	The plan and features match what staff use.
Help center and product docs	Zero data retention, retention period, audit log, residency scope, excluded features	The date you read it. These pages change often.
Trust center	SOC 2 report, subprocessor list, system card	The report period and the product it covers.
Your own files	Acceptable use policy, AI inventory, impact assessment	Each has an owner and a review date.

SUPERWISE® Sentinel, in these terms

Sentinel is an AI gateway that runs on the PC. Each statement below is from superwise.ai/sentinel/ or [/pricing/](https://superwise.ai/pricing/), read on October 8, 2026 [26], set against the terms in this kit.

TERM	WHAT SENTINEL DOES	ITS SCOPE
Guardrail	Replaces emails, phone numbers, account and card numbers, government IDs, and API keys with placeholders such as <EMAIL> before the request leaves.	Requests that pass through it. Names, street addresses, and passwords are absent from the published list.
Audit log	Every inspected interaction lands in an Activity Log: what was checked, what was redacted, and when.	Interactions Sentinel inspects. Traffic that never reaches it leaves no entry.
Data residency	It runs on the machine or in your own cloud, and nothing is routed through SUPERWISE.	Prompts go from the PC to the AI provider you already use, under that provider's terms.
AI gateway	Covers Claude Code, and claude.ai and chatgpt.com through a browser extension in preview for Chrome and Edge on Windows. Other chat sites are listed as in development.	The traffic sent through it. It finds no AI tools on its own, so it is no answer to shadow AI by itself.

- **No compliance claim.** Sentinel supplies no business associate agreement and makes no tool HIPAA compliant. It is one control inside your own program.
- **Where it runs.** Windows 10 and 11, with Claude Code also supported on Mac and Linux.
- **What it costs.** Free for 30 days with no card, then \$10, \$25, or \$50 a month per deployment.

Ask us the same questions. Every question in parts 2 to 4 applies to SUPERWISE too: which traffic passes through Sentinel, what the Activity Log records, and how long you keep it. Details are at superwise.ai/sentinel/.

Sources

Vendor pages were read on the date shown. Several show only a relative update date; the one shown on the day is recorded with each.

- 1 OpenAI, *Data controls in the OpenAI platform*. Zero Data Retention approval, abuse monitoring retention, and endpoint eligibility. Read October 8, 2026. developers.openai.com/api/docs/guides/your-data
- 2 Anthropic Privacy Center, *I have a zero data retention agreement with Anthropic. What products does it apply to?* Marked "Updated this week," read October 8, 2026. privacy.claude.com/en/articles/8956058
- 3 Anthropic Privacy Center, *Data retention practices for covered models*, dated September 5, 2026, read October 8, 2026. privacy.claude.com/en/articles/15425996
- 4 OpenAI, *Enterprise privacy at OpenAI*. Read October 8, 2026. openai.com/enterprise-privacy/
- 5 Anthropic Privacy Center, *Is my data used for model training?* Updated August 18, 2026, read October 8, 2026. privacy.claude.com/en/articles/7996868
- 6 OpenAI Help Center, *Data residency and inference residency for ChatGPT*. Marked "Updated 4 days ago," read October 8, 2026. help.openai.com/en/articles/9903489
- 7 Google Workspace Admin Help, *Gemini for Google Workspace FAQ*. Read October 8, 2026. knowledge.workspace.google.com/admin/gemini/gemini-for-google-workspace-faq
- 8 HHS Office for Civil Rights, *Guidance Regarding Methods for De-identification of Protected Health Information*. Read October 7, 2026. hhs.gov/hipaa/for-professionals/special-topics/de-identification/
- 9 HHS Office for Civil Rights, *Business Associates guidance*. Read October 6, 2026. hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/
- 10 OpenAI Help Center, *HIPAA eligible products and functionality*. Read October 7, 2026. help.openai.com/en/articles/20001069
- 11 Anthropic Privacy Center, *Business Associate Agreements (BAA) for Commercial Customers*. Read October 7, 2026. privacy.claude.com/en/articles/8114513
- 12 Regulation (EU) 2016/679, the General Data Protection Regulation, Article 4 (definitions of controller and processor) and Article 28 (processor contracts and subprocessors). eur-lex.europa.eu/eli/reg/2016/679/oj
- 13 HHS, *Breach Notification Rule*. Read October 7, 2026. hhs.gov/hipaa/for-professionals/breach-notification/
- 14 AICPA & CIMA, *System and Organization Controls: SOC suite of services*, and its publication on SOC 2 examinations. Read October 8, 2026. The landing page names the five categories; the report period and system description are features of the report itself. aicpa-cima.com/resources/landing/system-and-organization-controls-soc-suite-of-services
- 15 NIST Computer Security Resource Center glossary, *prompt injection*, from NIST AI 100-2e2025. Read October 8, 2026. csrc.nist.gov/glossary/term/prompt_injection
- 16 OWASP GenAI Security Project, *LLM01:2025 Prompt Injection*. Direct and indirect injection. Read October 8, 2026. genai.owasp.org/llmrisk/llm01-prompt-injection/
- 17 National Institute of Standards and Technology (July 26, 2024), *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1. doi.org/10.6028/NIST.AI.600-1
- 18 SUPERWISE, *AI glossary: the terms that carry weight*, entries for AI gateway and LLM gateway. Our own page, reviewed October 7, 2026. superwise.ai/learn/glossary/
- 19 National Institute of Standards and Technology (September 2020), *Security and Privacy Controls for Information Systems and Organizations*, SP 800-53 Rev. 5, including the audit and accountability (AU) controls. doi.org/10.6028/NIST.SP.800-53r5
- 20 Microsoft Learn, *Understand Shadow AI in Microsoft 365 admin center*. Describes a preview feature; updated August 2026, read October 8, 2026. learn.microsoft.com/en-us/microsoft-365/admin/manage/agent-shadow-ai
- 21 Regulation (EU) 2024/1689, the EU Artificial Intelligence Act, Article 14 (human oversight). eur-lex.europa.eu/eli/reg/2024/1689/oj
- 22 SUPERWISE Learn kit, *AI acceptable use policy template*, and the free template at superwise.ai/learn/ai-acceptable-use-policy/. Our own material.
- 23 National Institute of Standards and Technology (January 26, 2023), *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, the Govern and Map functions. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
- 24 OpenAI Help Center, *OpenAI Compliance Platform for Enterprise and Edu customers*. Marked "Updated 11 hours ago," read October 8, 2026. help.openai.com/en/articles/9261474
- 25 Mitchell, M., et al. (2019), *Model Cards for Model Reporting*, Proceedings of FAT* 2019. doi.org/10.1145/3287560.3287596
- 26 SUPERWISE, *Sentinel and Pricing*. Every product statement in part 6. Read October 8, 2026. superwise.ai/sentinel/ · superwise.ai/pricing/

The system card entry cites no source: the term has no standard definition, and vendors use it differently. Definitions here are plain-English summaries for business owners, checked on October 6 to 8, 2026. Check the source before you rely on one.